UNIVERSITY OF MARYLAND AT COLLEGE PARK

INSTITUTE FOR ADVANCED COMPUTER STUDIES

FINAL REPORT

Contract NCC 2-560
National Fellowship Program in
Parallel Processing
September 1, 1988 - December 31, 1992

Dr. Larry Davis, Principal Investigator
University of Maryland Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742

Final Report

1. Introduction - Summary. For the past three years the University of Maryland Institute for Advanced Computer Studies has administered a program of support for American graduate students who are conducting experimental research in parallel processing. The program has supported nearly 30 students during this time, many of whom have completed their Ph. D.s and taken prestigious jobs in academia, government and industry. These students were chosen from a large pool of applicants after a thorough review of written proposals by some of the nation's leading experts in the field of parallel computing.

The students were brought together once each year in a workshop where they presented the results of their research and interacted with senior researchers. These workshops were held in conjunction with major conferences on parallel processing, giving the program more publicity and providing the students an even greater opportunity to discuss their research and the program. Students were placed in government laboratories and industry during the summer to encourage them to pursue career opportunities that would subsequently involve them in government programs.

Our programmatic goals were to:

1) **Broaden the base of research support for parallel processing in the nation.** Our students were drawn from many disciplines and many different Universities. The program enabled us both to identify pockets of excellence in parallel computing, which might not in themselves justify funding but which are nonetheless of significantly high quality, and to amplify funding of especially important and high quality research programs.

2) **Identify the potential leaders in future generations of scientists and engineers.** By identifying the very best students in parallel processing, and bringing them together with each other and today's leaders in the field, we helped to lay the foundation for the emergence of the next generation of leaders in the field.

3) **Encourage high quality researchers to take positions in Government Laboratories and industry after completing their degrees.** Our yearly workshop brought the students into contact with representatives from government laboratories and industry. The summer internship program, put in place during 1991, placed the students at government contractors and laboratories.

While students from important application areas, both within and outside of computer science, were supported by the program, most of the students studied central issues in systems and language support for parallel processing.

2. Structure of the program

- **Who was eligible?**

The program provided support only to American citizens. Students must have had completed all course and written examination requirements for their Ph. D.s, and must have had a well defined dissertation research topic. The program placed a strong emphasis on **experimental** research in parallel computing and did not, in general, provide support to students whose research was essentially theoretical. There was a limit of one fellow in any single academic department,

although students from different academic departments in the same University could receive support.

With these restrictions our funds were used to support research rather than general graduate education and the students, being American citizens, were most likely to be future contributors to Department of Defense programs.

- **How were students chosen?**

There was one call for proposals each year during the late Fall, with proposals due early the following year. The program was advertised through direct mailings to all major computer science, electrical engineering and physics departments, announcements in publications like Communications of the ACM, IEEE Computer and SIG notices of the ACM, distribution of flyers describing the program at major conferences and workshops, as well as messages posted to interest groups on major computer networks. We received between 45-60 proposals in response to each Call for Proposals.

Proposals contained a brief (3-5 page) description of the research already performed and planned by the student, copies of vitae for both the student and his/her advisor as well as a letter of recommendation from the student's advisor.

Each proposal was read by two members of the evaluation committee, comprised of leaders from government and universities. Occasionally, when a proposal was on a topic not within the joint areas of expertise of the evaluation committee, it was sent to a special outside reviewer for comments. Each reviewer ranked the proposals that he/she considered, and indicated for the very best proposals why they merited funding. In March, we held a meeting of a subgroup of the evaluation committee at which the fellows were chosen. The first two meetings were held at M.I.T., while the 3rd year's evaluation meeting was held at Sandia National Laboratories.

Many excellent researchers from university and government have contributed their time during the past three years to reviewing student proposals. Appendix 1 lists the reviewers for each year.

- **Yearly workshop**

Each year we organized and held a workshop in which all of the students were brought together for two days to present the results of their research and interact with leading researchers from government, industry and academia. Appendix 2 contains the lists of speakers for the workshops. The 1990 workshop was held in conjunction with the 3rd Frontiers of Massively Parallel Computing Conference in College Park Maryland, and all students were provided support to attend that conference as well as the workshop. The 1991 workshop was held in conjunction with Supercomputing '91 and provided the students with introductions to the research activities at both Sandia National Laboratories and Los Alomos National Laboratories.

These workshops have been highly successful for a variety of reasons;

a) They helped to establish a camaraderie amongst the fellows, and laid the foundation for cooperation among the next generation of leaders in the field of parallel computing.

b) They provided a valuable recruiting forum for representatives from government, academia and industry. We extended invitations to many representatives from government agencies, laboratories and manufacturers of parallel processors to attend the workshops and meet the students.

c) Especially for students who were seeking renewals of their fellowships, it provided us with an opportunity to review their work and judge its quality.

- **Internships**

In 1991 we began a program of placing fellows as summer interns in government laboratories, government supported supercomputing facilities, and companies at the leading edge of parallel processing. This extension to the program enhanced the likelihood that our fellows will choose careers that will directly benefit government programs. Students interned as follows:

1) Jeffrey Becker from UC Davis worked at NASA Ames.

2) Peter Beckman from Indiana University worked at BBN in Cambridge in John Goodhue's group.

3) Matt Blumrich from Princeton spent the summer at Intel with Paul Close working on graphic support for Touchstone.

4) Mark Crovella from the University of Rochester worked at Cornell's Supercomputing Center.

5) Diane Davison from the University of Colorado interned at Intel.

6) Matt Evett from the University of Maryland worked at Thinking Machines building AI systems on the CM.

7) Greg Frazier from UCLA spent time at Hughes Research Laboratory in Malibu.

8) Greg Huber, a physicist at Boston University, spent a summer at Los Alamos.

9) Joe Hummel from the University of California at Irvine worked at the Supercomputing Center at University of Illinois - Urbana.

10) Todd Jochem from CMU interned at Martin Marietta in Denver with Rick Lukas.

11) Raymond Loy of RPI worked at Thinking Machines.

12) Catherine McCann from the University of Washington interned at Terra Computer Group.

13) Brian Totty from the University of Illinois spent eight weeks at MIT.

14) Richard Zucker from the University of Washington was interested in cache coherency problems and spent a summer at Lawrence Livermore Laboratories.

3. History of the program

In choosing students for support we considered the following three factors:

1) Quality of the students and their research projects (the most important criteria for selection).

2) Intellectual balance of the program with research projects involving both both basic research in parallel processing and in applications of parallel computing that were deemed critical to government programs.

3) Equal opportunity considerations, to attract more women and minorities into computer science and engineering, generally, and parallel computing, specifically.

Appendix 3 contains a list of the students supported each year, indicating their institution, and research area. For those students who have finished their degrees, we also indicate their current employment, when available, in a footnote. In the following section we outline the research projects of each student. Most of the first group of students supported by the program have completed their degrees and are now working in industry, academia and government.

We were highly successful in identifying exceptionally qualified students for support under the program. Just a few highlights of the program and the accomplishments of the students who constituted our first group of fellows:

1) Ann Rogers from Cornell is now on the Computer Science faculty at Princeton University.

2) David Goldstein, an aeronautical engineer from CalTech, is now a postdoctoral scientist at Brown University.

3) Donovan Schneider from the University of Wisconsin won the best paper award at the 1990 ACM SIGMOD conference.

4) Greg Byrd, who completed his degree at Stanford is now working for Digital's Western Research Laboratory.

5) Charles Tong, a computational numerical analyst from UCLA, is a postdoctoral scientist at Sandia National Laboratories.

4. Research projects: 1988-1992. In this section we provide brief descriptions of the research conducted by the fellows under the program. We separate the projects into three categories: systems, languages, and scientific computation

4.1 Systems. Research in systems has focused mainly on problems of memory management, modeling of cache systems for shared memory machines, and design and analysis of I/O subsystems for shared memory machines. The specific projects include:

• Gregory Byrd, Stanford University (Advisor: Michael Flynn): Dr. Byrd, now with DEC Western Research Laboratories, studied problems related to message passing in shared memory multiprocessors. Streamlines are an implementation of streams in a shared memory system which allow cache-to-cache transfers and synchronization for interprocess messages. He defined a set of behavioral models for streamline-based message passing architectures, coded benchmarks including a PDE solver and a VLSI floor plan optimizer, and performed a variety of benchmarking experiments.

• Brian Marsh, Rochester University (Advisor: Tom LeBlanc): Mr. Marsh studied the role of virtual memory in large scale NUMA (Non-uniform memory access) machines. His experimental research was conducted on the BBN Butterfly. A key concern in Mr. Marsh's research was the ability to scale the memory management model to machine size. His memory management hierarchy contains four levels - local memory, global memory, backing store and user-level view of virtual memory. All machine dependent memory management is isolated in the lowest level of the hierarchy, and by implementing this on a per-processor basis, he believes that his memory management mechanisms would sale up to much larger multiprocessor systems. Mr. Marsh has published the results of his research extensively, with over seven articles in major national conferences and workshops.

• William Bolosky, Rochester University (Advisor: Tom LeBlanc): Mr. Bolosky studied automatic locality management in parallel programming systems. He has designed and implemented a memory management system that automatically migrates and replicates memory pages in Mach. His overriding research goal was to determine the most cost-effective combination of hardware and software support for the management of data locality in large shared-memory multiprocessors. His experimental research was performed on a BBN Butterfly machine.

• Brian Totty, University of Illinois (Advisor: Dan Reed): Mr. Totty has been studying system and software techniques to improve the performance of parallel programs executing on NUMA architectures. The aim of the research was to develop policies that improve system performance by providing data management at a fine level of granularity. Specifically, he was pushing the level of data placement control much deeper than is usually done, so that each algorithmic data structure can benefit from its own data management policy. In order to develop such data structure dependent management tools, he has implemented a tracing tool that allows him to compare different data management policies and provide statistical summary information to applications programmers. He is focusing on data structures that arise in scientific computing and has studied 8 algorithms including matrix multiplication, Cholesky decomposition and FFT algorithms. His research was related to the more language-oriented research conducted by Dr. Ann Rogers from Princeton during the first two years of the program.

• David Kotz, Duke University (Advisor: Carla Ellis): Dr. Kotz's research involved designing file systems for MIMD systems. Specifically, he compared, both theoretically and empirically, the effects of alternative prefetching strategies on the average block read time for files. Previous research on this problem involved knowing, a priori, the pattern of fetches to files. Dr. Kotz was able to show that using only weaker statistical models of fetching patterns one could obtain

comparable I/O performance. Dr. Kotz's research been described in a major paper in the IEEE Transactions on Parallel and Distributed Systems.

• Donovan Shneider, University of Wisconsin (Advisor: David DeWitt). Dr. Shneider developed parallel algorithms for fundamental database operations. He constructed a simulator for a large multiprocessor database machine and implemented and evaluated the performance of several parallel join operators and query optimization algorithms in the simulator. The algorithms were also implemented on a 32 processor Intel iPSC/2 with nearly linear speedup. Dr. Shneider received a best student paper award at the annual SIGMOD conference for his research.

• Jeffrey Becker, UC - Davis (Advisor: Arvin Park): Mr. Becker used trace driven simulation to analyze the I/O requirements of parallel processing systems. He focused on scientific computations - specifically certain signal processing and global change codes - to test his ideas. His research was conducted on a variety of parallel machines, including the Multimax and the Intel hypercube.

•Matthias Blumrich, Princeton (Advisor: Kai Li): Mr. Blumrich was interested in scientific visualization for computations performed on multicomputers and worked to enhance the Intel-DARPA Touchstone multicomputer to provide direct hardware support for visualization of dynamic 3-D processes. Mr. Blumrich spent part of the summer of 1991 at Intel.

• Gregory Frazier, UCLA (advisor: Yuval Tamir): Mr. Frazier's research was performed as part of the UCLA ComCoBB Project. The goal of that project was to design and implement a single chip high performance communication coprocessor for use in VLSI multicomputer systems. His specific research dealt with the design and implementation of the dynamically allocated, multique (DAMQ) buffer, the internal buffer of a small nxn VLSI communication switch. He was particularly interested in minimizing the possibly of tree saturation (the tree of nodes blocked by the development of a hot spot) in such networks. He investigated the effect of different buffer space management policies on the development of hot spots.

• Richard Zucker, University of Washington (Advisor: Jean-Loup Baer): Mr. Zucker developed methods to minimize performance penalties due to cache coherence effects in large shared memory multiprocessors. One of the principal causes of cache misses in such systems is memory latency, since it can take more than 100 cycles in a multistage interconnection network for a line to be returned from memory. He investigated how much performance improvement can be obtained by features such as write buffers, prefetching, and lock-up free caches. He conducted his experimental research on a Sequent Symmetry Model B. Mr. Zucker spent part of the summer of 1991 at Lawrence Livermore National Laboratory where they have developed an extensive simulation package for studying cache management problems in shared memory machines.

• David Lilja, University of Illinois (Advisor: Pen-Chung Yew): Mr. Lilja was also interested in cache coherence protocols for shared memory multiprocessors. His research goal was to examine the performance of several different latency reduction and hiding techniques to determine how they may be adapted for use in a multiprocessor system such as Cedar. Specifically, he studied the performance of caches versus prefetching and context switching.

• Eric Schwabe, M.I.T. (Advisor: Tom Leighton): Dr. Schwabe, who accepted a faculty position at Northwestern University, studied the ability of different interconnection networks to simulate one another efficiently. The major result of his dissertation was a demonstration that all common interconnection networks (e.g., shuffle exchange, butterfly, etc.) can efficiently simulate one another. He more recently has been looking at online algorithms for scheduling work in distributed computing systems.

• Mark Crovella, University of Rochester (Advisor: Thomas LeBlanc). The goal of Mr. Crovella's research was to integrate performance analysis, prediction and tuning into the parallel program development process so as to avoid extensive and error-prone program modifications. Mark published papers that discuss the role of performance prediction, using analytic models and rapid experimentation to select parallelizations, and a tool for rapid performance assessment based on performance predicates. During Mr. Crovella's internship at the Conell Supercomputing Center he studied methods of parallelizing combinatorial search algorithms on the Center's KSR machine.

• Katherine McCann, University of Washington (Advisor: John Zahorjan). Ms. McCann's research focuses on scheduling parallel programs on HPC's that are multiprogrammed. By supporting multiprogramming on these machines, we can provide greater availability of an expensive resource, better aggregate response time to a community of users and increased utilization of the facility. Her research has focused on two specific problems: one, in which she assumes that a single job may be executed on any number of processors, even one, and a second in which jobs have minimum memory requirements to run. For each case, she has developed policies for scheduling. Ms. Davisson conducted her internship at Tera Computer Company in Seattle, working on their scheduling system.

• Erich Nahum, University of Massachusetts (Advisor: Don Towsley). Mr. Nahum has been studying problems related to experimentally evaluating approaches for exploiting parallelism in system support for high speed networks. He has focused on reducing the bottleneck in network performance posed by crossing the user-kernel boundary, and is showing that parallelism can be used to overcome this bottleneck. His research is conducted using ARPA sponsored platforms, including the Mach operating system and the x-Kernal. He spent his internship at the University of Arizona, working on the design and implementation of a MP safe version of the x-Kernel.

4.2 Languages. Research in languages has concentrated mainly on problems relating to concurrency specification and control, and automatic parallelization of codes for scientific computation.

• Anne Rogers, Cornell University (Advisor: Keshav Pingali): Dr. Rogers's research focused on the exploitation of spatial locality - the placement of data - to reduce the effective latency of memory reads in multiprocessors. Unlike the research of Brian Totty, which focuses on tools for supporting the application programmer in distributed data in NUMA machines, Dr. Rogers was more interested in providing additional functionality to parallelizing compilers for automatic distribution of data. Dr. Rogers implemented a hydrodynamics benchmark from Los Alamos consisting of 100 lines of FORTRAN. For a 32 processor Intel iPCS/2 the code generated by her compiler achieved a speedup of 17.4 . A hand-crafted version of the same program achieved a speed up of about 26. Dr. Rogers is now a member of the Computer Science Department at Princeton University.

• Mark Hansen, M.I.T. (Advisor: Frank Leighton): Dr. Hanson's research was concerned with applications of computational complexity to parallel processing problems. He was specifically concerned with problems of embedding graphs in the plane. These algorithms have important applications in communication load balancing, dynamic allocation of jobs to processors, reconfiguring around faults in interconnection networks, and the dynamic configuration of distributed computing networks.

• Carolyn Craig Williams, University of Virginia (Advisor: Paul Reynolds). Ms. Williams research involved developing mechanisms for concurrency control in asynchronous parallel

computations. Most concurrency control mechanisms are based on locks and delays. Ms. Williams investigated alternative concurrency control mechanisms based on a system of locally synchronized logical clocks. She introduced a new synchronization primitive - the **parop** - and has developed methods for implementing parops without locking or the risk of rollback. She studied the use of parops to form the basis for an efficient, scalable solution to the cache coherency problem.

• Kathryn McKinley, Rice University (Advisor: Ken Kennedy): Ms. McKinley's research sought to determine the potential success of automatic parallelizers for shared memory multiprocessors, and to determine if the types of analysis and techniques needed for good code generation are currently available or can be automated. Her research focused on a specific set of benchmarks - RICEPS (the Rice test suite) and the Perfect Benchmarks. Her research was conducted on both a Sequent Symmetry and a BBN Butterfly.

• Peter Newton, University of Texas (Advisor: James Browne): Mr. Newton's research was conducted in the context of the CODE/ROPE parallel programming research environment project at Texas. His research involved laying the foundations for efficient execution of a high level specification language for parallel programs. The research began with the development of a model of execution which defines a target language for the high level specification language of CODE and ROPE and can be translated to a spectrum of shared-name space and partitioned name-space architectures.His target languages were Ada on a Sequent shared memory multiprocessor and C on the SYMULT.

• Matt Evett, University of Maryland (Advisor: Jim Hendler): Dr. Evett was interested in the application of massive parallelism to problems in AI. His research focused on two specific problems. First, he developed a parallel knowledge representation language called PARKA that supports rather general inheritance hierarchies. Second, he has developed a parallel heuristic search algorithm that shows the best speedup of any parallel search algorithm to date. Dr. Evett spent the summer of 1991 with Thinking Machines.

• Lorenz Huelsbergen, University of Wisconsin (Advisor: James Laurus). Mr. Huelsbergen studied hybrid parallelization techniques composed of static and dynamic components. Fast static analyses provide partial side effect information during compilation. Other information, difficult to obtain at compile time, is determined at run time. His research is especially applicable to problems involving complex and irregular data structures and to analysis of programs written in weakly typed languages such as LISP. His research was extended so that it could be applied to programs written in FORTRAN or C++.

• Joseph Hummel, UC Irvine (Advisor: Alexandru Nicolau). Mr. Hummel conducted research on analyzing and transforming codes using pointer analysis and he created a general dependence test for pointers. He also worked on implementing his pointer analysis technique in an ANSI C compiler, creating a pointer-based benchmark suite and supporting debugging of pointer data structures. This implementation allows for the collection of results on the effectiveness of the approach, as well as a demonstration of its viability. Many of the recently identified national challenge areas which will require large scale parallel manipulation of complex data structures based on trees and graphs will require techniques such as those developed by Mr. Hummel, for their successful implementation on HPC's.

4.3 Scientific computing and applications. We have supported several outstanding students interested in either fundamental problems in scientific computing or in important applications of parallel processing to the physical sciences. Most of the research in this area has

focused on the use of massive parallelism to solve problems in scientific computing.

• David Goldstein, California Institute of Technology (Advisor: Bruce Sturtevent). Dr. Goldstein received his degree in Aeronautical Engineering. His research involved the use of parallel processing to study molecular dynamics. He developed efficient methods for calculating fluid flows far from equilibrium and at modest Reynolds number. This research has extended the understanding of complex gas dynamics systems made up of a large number of simple molecular elements.His research was conducted on a 192 processor Symult 2010 multicomputer. Dr. Goldstein is currently a postdoctoral scientist at Brown University.

• Richard Strilka, Boston University (Advisor: Claudio Rebbi). Dr. Strilka worked with Nobel prize winning physicist Caludio Rebbi on simulating flux vortices using the CM. Being able to model the dynamics of such vortices is critical to understanding new high temperature superconductors. One way to construct a high temperature superconductor is to evaporate the superconducting ceramic onto a substrate. The ceramic crystal, however, does not form perfectly; small islands of perfect crystals form in clusters separated by boundaries. Dr. Strilka has generalized the two dimensional modeling techniques developed by Chris Lobb at Harvard so that he can study the properties of these superconductors in three dimensions.

• Charles Tong, UCLA (Advisor: Tony Chan): Dr. Tong studied efficient preconditioners for the preconditioned conjugate gradient method for solving elliptic PDE's. He developed a class of multilevel preconditioners based on concepts from digital signal processing, called multilevel filtering preconditioners. These preconditioners were implemented on some second order elliptic PDE's including anisotropic problems. His experimental research was conducted using both an Alliant FX/80 and a CM. Dr. Tong's dissertation research resulted in 3 major journal publications. He is currently a research scientist at Sandia National Laboratories.

• Peter Beckman, Indiana University (Advisor:David Weiss): Mr. Beckman studied the use of quadtree representation of sparse matrix to solve, in parallel, large, sparse linear systems of equations. He developed a system of programs that allow him to experiment with pivot strategies, error fill-in, and selection of non-singular blocks as pivot candidates for Jordan elimination.

• Raymond Loy, RPI (Advisor: Joseph Flaherty): Mr. Loy's research is on the subject of efficient parallel solutions of the three-dimensional Euler equations of compressible fluid dynamics. Such problems are generally regarded as being amongst the most demanding in scientific computations. Mr. Loy is also studying the used of variable resolution techniques to study problems in scientific computation. Since his problems are three-dimensional, he is employing techniques based on octtrees. His research is being conducted on a MASPAR available at RPI.

• Greg Huber, Boston University (Advisor: H. Eugene Stanley): Mr. Huber's research involved using the Connection Machine to pursue a new program concerning novel physics that occurs in weakly turbulent systems. This work has great promise for altering, in a fundamental way, the fashion in which we understand pattern formation in chemical oscillations and reaction diffusion systems. Mr. Huber spent the two summers working at Los Alomos National Laboratories.

• Ken Cox, Washington University (Advisor:Gruia-Catalin Roman): Mr. Cox has studied methods for visualizing the activity of programs in multiprocessor systems. He has developed a visualization package that allows applications programmers to associate graphical objects with program objects and define the relationship between changes in their images and their values.

• Diane Davisson, University of Colorado (Advisor: Goetz Graefe). Ms. Davisson studied how to exploit parallelism including I/O capabilities for database query execution by designing and

validating a scalable, decentralized algorithm to optimally manage multiple resources. She built and experimented with a simulator that led to new methods of resource allocation for query processing in highly parallel database systems. She has applied analysis techniques from microeconomics to these problems.

• Todd Jochem, CMU (Advisor Charles Thorpe). Mr. Jochem has been studying novel neural network architectures for complex problems in visual navigation. His research directly supports ARPA's Unmanned Ground Vehicle program, which requires fast and robust navigation of military vehicles over roads and cross country. Mr. Jochem has focused on the problem of integrating neural architectures, each optimized to solve a specific navigation problem into a complete system that cna operate in a wide spectrum of conditions. He has spent his internship at Martin Marietta, the integrating contractor on the UGV program, porting his neural network systems to their vehicle. His implementation work is done on the iWARP machine, a systolic array processor designed at CMU and manufactured by Intel.

APPENDIX 1

Evaluation Committee
1989

Dr. William Dally
MIT
Cambridge, MA 02139

Dr. Jack Dongarra
University of Tennessee
Knoxville, TN 37996

Dr. Danny Hillis
Thinking Machine Corp.
Cambridge, MA 02142

Prof. Joseph Ja'Ja'
University of Maryland
College Park, MD 20743

Prof. Tom Leighton
MIT
Cambridge, MA 02139

Evaluation Committee
1990

Dr. William Dally
MIT
Cambridge, MA 02139

Dr. Jack Dongarra
University of Tennessee
Knoxville, TN 37996

Dr. Joseph Ja'Ja'
University of Maryland
College Park, MD 20743

Dr. Tom LeBlanc
University of Rochester
Rochester, NY 14627

Dr. Tom Leighton
MIT
Cambridge, MA 02139

Evaluation Committee
1991

Prof. Tony Chan
University of CA, Los Angeles
Los Angeles, CA 90024

Dr. William Dally
MIT
Cambridge, MA 02139

Dr. Sudip Dosanjh
Sandia National Laboratories
Albuquerque, NM 87185-5800

Prof. Carla Ellis
Duke University
Durham, NC 27706

Prof. Daniel Reed
University of Illinois
Urbana, IL 61801-2987

Evaluation Committee
1992

Dr. Richard Allen
Sandia National Laboratory
Albuquerque, NM 87185-5800

Prof. Jack Dongarra
University of Tennessee
Knoxville, TN 37996

Prof. Ken Kennedy
Rice University
Houston, TX 77251

Prof. Daniel Reed
University of Illinois
Urbana, IL 61801-2987

Prof. Pete Stewart
University of Maryland
College Park, MD 20740

Prof. Bill Wulf
University of Virginia
Charlottesville, VA 22903

Prof. Rama Chellapa
University of Maryland
College Park, MD 20740

Prof. Joseph Ja'Ja'
University of Maryland
College Park, MD 20740

APPENDIX 2

Workshop
on
Parallel Processing

Monday, August 28, 1989

8:30 - 9:00	Continental Breakfast
9:00 - 9:10	Larry Davis, U of MD Institute for Advanced Computer Studies *Introduction*
9:10 - 9:55	Mark Pullen, Defense Advanced Research Projects Agency *DARPA Research Programs in High Performance Computing*
9:55 - 10:40	Jim Browne, University of Texas *The Software Engineering of Parallel Programs*
10:40 - 10:55	Coffee Break
10:55 - 11:30	David Goldstein, California Institute of Technology *A Parallel Method for Computing the Flow of Discrete Velocity Gases*
11:30 - 12:00	Donovan Schneider, University of Wisconsin *Complex Query Processing in a Multiple Processor Database Machine*
12:00 - 1:30	Lunch - Saturn & Venus Room
1:30 - 2:15	Tom Leighton, Massachusetts Institute of Technology *Recent Results on Packet Routing Algorithms for Fixed-Connection Networks*
2:15 - 2:30	Charles Tong, University of California - Los Angeles *Parallel Iterative Methods for Elliptic PDE's and Their Implementation*
2:30 - 2:45	Richard Strilka, Boston University *A Massively Parallel Simulation of Flux Vortices in Superconductors and Cosmic Strings*
2:45 - 3:00	Brian Marsh, University of Rochester *The Integration of Policy and Mechanism in NUMA Virtual Memory*
3:00 -3:15	Refreshments
3:15 - 4:00	Jack Dongarra, Argonne National Laboratory *LAPACK: A Linear Algebra Library for High Performance Computers*
4:00 - 4:45	Ken Kennedy, Rice University *Programming Support Environments for Parallel Computer Systems*
5:00 - 7:00	Smithsonian National Air and Space Museum
7:30	Dinner

Workshop
on
Parallel Processing

Tuesday, August 29, 1989

8:30 - 9:00	Continental Breakfast
9:00 - 9:45	Paul Schneck , Supercomputing Research Center *A Reconfigurable Systolic Array*
9:45 - 10:30	Bill Wulf, National Science Foundation *The WM Computer Architecture*
10:30 - 10:50	Coffee Break
10:50 - 11:20	Anne Rogers, Cornell University *Process Decomposition Through Locality of Reference*
11:20 - 11:35	David Kotz, Duke University *High Performance File System Design for MIMD Parallel Processors*
11:35 - 12:15	Azriel Rosenfeld, University of Maryland *Parallel Vision*
12:15 - 1:30	Lunch - Saturn & Venus Room
1:30 - 2:15	William Dally, Massachusetts Institute of Technology *Universal Mechanisms for Concurrency*
2:15 - 2:30	Carol Craig Williams, University of Virginia *Parallel Operations on Value Sequences*
2:30 - 2:45	Mark Hansen, Massachusetts Institute of Technology *Approximation Algorithms for Geometric Embedding in the Plane with Applications to Parallel Processing Problems*
2:45 - 3:00	Gregory Byrd, Stanford University *Support for Message Passing in Shared Memory Multiprocessors*
3:00 - 3:15	Refreshments
3:15 - 4:00	Joseph Ja'Ja', University of Maryland *Parallel Algorithms for VLSI Routing*

END OF WORKSHOP

Workshop on Parallel Processing
University of Maryland Center of Adult Education
October 10 - 11, 1990

Wednesday, October 10

1:00 - 1:05	Welcome - Larry Davis *Director, Univ. of MD Institute for Advanced Computer Studies*
1:05 - 1:50	Tom LeBlanc *University of Rochester*
1:50 - 2:35	Bill Dally *Massachusetts Institute of Technology*
2:35 - 2:55	Peter Beckman *Indiana University*
2:55 - 3:15	Kenneth Cox *Washington University - St. Louis*
3:15 - 3:55	Michael Duff *University College - London*

(4:00 - 6:00 - Frontiers '90 Panel Discussion)

7:00	DINNER Far East Restaurant

Thursday, October 11

8:30 - 9:00	Continental Breakfast
9:00 - 9:45	H.J. Siegel *Purdue University*
9:45 - 10:30	Carla Ellis *Duke University*
10:30 - 10:50	Matthew Evett *University of Maryland*
10:50 - 11:10	Break
11:10 - 11:30	Brian Marsh *University of Rochester*
11:30 - 11:50	Eric Schwabe *Massachusetts Institute of Technology*
11:50 - 12:30	Clyde Kruskal *University of Maryland*
12:30 - 1:30	LUNCH
1:30 - 1:50	Richard Strilka *Boston University*
1:50 - 2:10	Brian Totty *University of Illinois*
2:10 - 2:30	Carolyn Williams *University of Virginia*
2:30 - 3:15	Paul Schneck *Supercomputing Research Center*
3:15 - 3:35	David Kotz *Duke University*
3:35 - 3:55	David Lilja *University of Illinois*
3:55 - 4:10	Closing Remarks

Workshop on Parallel Processing
Doubletree Hotel
November 18-19, 1991

AGENDA

Monday, November 18th

8:30 - 9:10	Continental Breakfast
9:10 - 9:20	Opening Remarks Prof. Larry S. Davis, University of Maryland Dr. Gil Weigand, DARPA
9:20 - 9:50	Dr. Stephen Squires, DARPA
9:50 - 10:15	Matt Blumrich, Princeton University *"An Empirical Comparison of Loop Scheduling Methods on a Multiprocessor"*
10:15 - 10:30	Coffee Break
10:30 - 10:55	Brian Totty, University of Illinois *"Experimental Analysis of Data Management for Distributed Data Structures "*
10:55 - 11:20	Kathryn McKinley, Rice University *"Evaluating the Potential for Automatic Parallel Code Generation"*
11:20 - 11:50	Dr. Dan Reed, University of Illinois *"Performance Environments for Scalable Parallel Systems"*
11:50 - 12:15	William Bolosky, University of Rochester *"Automatic Locality Management in Parallel Programming Systems"*
12:15 - 1:45	Lunch
1:45 - 2:15	Dr. Richard Allen, Sandia National Laboratory *"Sandia's Research in Massively Parallel Computing"*
2:15 - 2:40	Matthew Evett, University of Maryland *"Massively Parallel Heuristic Search Within Limited Memory"*
2:40 - 3:00	Coffee Break
3:00 - 3:25	Gregory Frazier, University of California at Los Angeles *"High Performance Buffered Communication Switches"*
3:25 - 3:55	Dr. Jean-Loup Baer, University of Washington *"Memory Management In Shared-Memory Processors"*
3:55 - 4:20	Jeff Becker, University of California at Davis *"An Analysis of the Information Content of Address Reference Streams"*

4:20 - 4:50	Karen Warren, Lawrence Livermore National Laboratory *"Parallel Fortran Preprocessor"*
6:00	Dinner
7:00 - 9:00	Supercomputing '91 Reception - La Posada

Tuesday, November 19th

8:30 - 9:15	Continental Breakfast
9:15 - 9:45	Dr. Mani Chandy, California Institute of Technology
9:45 - 10:10	Greg Huber, Boston University *"Pattern Formation, Turbulence and Coupled Maps"*
10:10 - 10:40	Dr. Geoffrey Fox, Syracuse University *"Parallel Computing Applications in the Real World"*
10:40 - 10:55	Coffee Break
10:55 - 11:20	Richard Zucker, University of Washington *"A Performance Study of Memory Consistency Models"*
11:20 - 11:50	Dr. Jill Mesirov, Thinking Machines Corporation *"Parallel N-body Algorithms"*
11:50 - 12:15	Peter Newton, University of Texas at Austin *"Translation of an Abstract Parallel Programming Model to Efficient Execution Structures Across a Variety of Target System Architectures"*
12:15 - 1:30	Lunch
1:30 - 2:00	Dr. Tony Chan, University of California at Los Angeles *"Hierarchical Algorithms in Parallel Scientific Computing"*
2:00 - 2:25	Peter Beckman, Indiana University *"Quadtree Algorithms for Sparse Matrices on the Butterfly"*
2:25 - 2:55	Dr. Jack Dongarra, University of Tennessee
2:55 - 3:20	Raymond Loy, Rensselaer Polytechnic Institute *"Parallel Adaptive Strategies for Three-Dimensional Compressible Flows"*
3:20 - 3:40	Closing Remarks

Workshop on Parallel Processing
Radisson Plaza Hotel
November 16-17, 1992

AGENDA

Monday, November 16th

8:30 - 9:00	Continental Breakfast
9:00 - 9:05	Opening Remarks Prof. Larry S. Davis, University of Maryland
9:05 - 9:35	Dr. Monica Lam, Stanford University *"Lingusitics and Compiler Support for Parallel Computing"*
9:35 - 10:00	Matt Blumrich, Princeton University *"SHRIMP:Scalable High-Performance Really Inexpensive Multiprocessor"*
10:00 - 10:25	Catherine McCann, University of Washington *"Scheduling Multiprogrammed Distributed Memory Multiprocessors"*
10:25-10:40	Coffee Break
10:40 - 11:05	Mark Crovella, University of Rochester *"Exploiting Early Performance Prediction in Parallel Programming"*
11:05 - 11:30	Diane Davison, University of Colorado *"Dynamic Resource Allocation for Parallel Database Query Execution"*
11:30 - 12:00	Dr. David Forslund, Los Alamos National Laboratory *"Progress in High Performance Distributed Computing"*
12:00 - 12:30	Gregory Frazier, University of California, Los Angeles *"Distributed Flow Control in Multicomputer Communication Networks"*
12:30 - 1:45	Lunch
1:45 - 2:15	Dr. George Sell, AHPCRC *"Fast Poisson Solver"*
2:15 - 2:40	Greg Huber, Boston University *"Vortex Decay"*
2:40 - 3:00	Coffee Break
3:00 - 3:25	Dr. Anat Agarwal, MIT
3:25 - 3:50	Lorenz Huelsbergen, University of Wisconsin-Madison *"Dynamic Parallelization of MOSTLY FUnctinal Languages"*
3:50 - 4:20	Dr. Burton Smith, Terra Computing *"Computing in Space and Time"*

4:20 - 4:45	Erich Nahum *"Exploiting Parallelism in High-Speed Networking"*
6:30	Supercomputing 92 Reception

Tuesday, November 17th

8:30 - 9:00	Continental Breakfast
9:00 - 9:30	Dr. Jim Schwarzmeier, Cray Research Inc. *"Research Opportunities at Cray Research"*
9:30- 9:55	Dr. Joel Saltz, University of Maryland *"Runtime Compilation for Scalable Multiprocessors"*
9:55 - 10:25	Joseph Hummel, University of California, Irvine *"Improving Pointer Analysis Through Abstract Data Structure Descriptions"*
10:25- 10:40	Coffee Break
10:40 - 11:05	Todd Jochem, Carnegie Mellon University *"MANIAC: A Next Generatin Neurally Based Autonomous Road Follower"*
11:05 - 11:25	Raymond Loy, Rensselaer Polytechnic University *"Parallel Adaptive Stratgies for 3 Dimensional Compressible Flows"*
11:25 11:55	Dr. Donald Truhlar, NSF Supercomputing Center *"Quantum Dynamics on Parallel Computers"*
11:55 - 12:20	Jeff Becker, University of California, Davis *"Exploiting Redundancy to Increase I/O PIN Brand Width"*
12:20 - 12:30	Closing Remarks
12:30 - 1:30	Lunch

APPENDIX 3

NASA ASSISTANTSHIP RECIPIENTS
1989

Gregory Byrd - Stanford University, Dept of Computer Science
Support for Message Passing in Shared Memory Multiprocessors

David Goldstein - California Institute of Technology
A Parallel Method for Computing the Flow of Discrete Velocity Gases

Mark Hansen - Massachusettes Institute of Technology, Dept of Mathematics
Approximation Algortihms for Geometric Imbedding in the Plane with Applications to Parallel Processing Problems

David Kotz -Duke University, Dept of Computer Science
High Performance File System Design for MIMD Parallel Processors

Brian Marsh - University of Rochester, Dept of Computer Science
The Integration of Policy and Mechanism in NUMA Virtual Memory

Anne Rogers - Cornell University, Dept of Computer Science
Process Decompostion Through Locality of Reference

Donovan Schneider - Unversity of Wisconsin, Dept of Computer Science
Complex Query Processes in a Multiple Processor Database Machine

Richard Strilka - Boston University, Dept of Physics
A Massively Parallel Simulation of Flux Vortices in Superconductors and Cosmic Strings

Charles Tong - University of California, Los Angeles, Dept of Mathematics
Parallel Iteractive Methods for Elliptic PDE's and Their Implications

Carolyn Craig Williams - University of Virginia, Dept of Computer Science
Parallel Operations on Value Sequences

NASA ASSISTANTSHIP RECIPIENTS
1990

Peter Beckman - Indiana University, Dept of Computer Science
Parallel Algorithms

Kenneth Cox - Washington University-St. Louis, Dept of Computer Science
Visualization of Concurrent Computations

Matthew Evett, University of Maryland, Dept of Computer Science
PARKA-PARallel Knowledge represetation and Association

David Kotz -Duke University, Dept of Comuter Science
High Performance File System Design for MIMD Parallel Processors

David Lilja - University of Illinois, Dept of Electrical Engineering
Memory Latency Reduction in Shared Memory Multiprocessors

Brian Marsh - University of Rochester, Dept of Computer Science
The Integration of Policy and Mechanism in NUMA Virtual Memory

Eric Schwabe - MIT, Dept of Computer Science
Parallel Computation on Hyercube-Derived Networks

Richard Strilka - Boston University
Flux Vertices in Superconductors

Brian Totty - University of Illinois, Dept of Computer Science
Experimental Analysis of Distributed Data Structures and Data Management

Carolyn Craig Williams - University of Virginia, Dept of Computer Science
Parallel Operations on Value Sequences

NASA ASSISTANTSHIP RECIPIENTS
1991-1992

Jeffrey Becker - University of California, Davis, Dept of EE & Computer Science
I/O Support for Parallel Processing

Peter Beckman - Indiana University, Dept of Computer Science
Parallel Algorithms

Matthias Blumrich - Princeton University, Dept of Computer Science
Hardware Design for Real Time Image Processing

William Bolosky - University of Rochester, Dept of Computer Science
Automatic Locality Management in Parallel Programming Systems

Matthew Evett - University of Maryland, Dept of Computer Science
PARKA-PARallel Knowledge represetation and Association

Gregory Frazier - University of Californis, LA, Dept of Comuter Science
High Performance Buffered COmmunication Switches

Greg Huber - Boston University, Dept of Physics
ttern Formatin in Spatio-Temporal Intermittency

Raymond Loy - Rensselaer Polytechnic Institute, Dept of Computer Science
Parallel Adaptive Strategies for Three-Dimensional Compressible Flows

Kathryn McKinley - Rice University, Dept of Computer Science
Evaluating the Potential for Automatic Parallel Code Generation

Peter Newton - University of Texas, Austin, Dept of Computer Science
Translation of an Abstract Parallel Programming Model to Efficient Execution Structures Across a Variety of Target System Architectures

Brian Totty - University of Illinois, Dept of Computer Science
Experimental Analysis of Distributed Data Structures and Data Management

Richard Zucker - University of Washington, Dept of EE and Computer Science
Synchronization Patterns in Parallel Processing

NASA ASSISTANTSHIP RECIPIENTS
1992-1993

Jeffrey Becker, University of California, Davis, Dept.of EE & Computer Science
I/O Support for Parallel Processing

Matthias Blumrich, Princeton Universtiy, Dept. of Computer Science
Hardware Design for Real Time Image Processing

Mark Crovella, University of Rochester, Dept. of Computer Science
Analyzing, Predicting & Tuning Parallel Program Performance

Diane Davison, University of Colorado, Computer Science Dept.
Dynamic Resource Allocation for Parallel Database Query Execution

Gregory Frazier, University of California, L.A., Dept. of Computer Science
High Performance Buffered Communication Switches

Greg Huber, Boston University, Dept. of Physics
Pattern Formation in Spatio-Temporal Intermittency

Lorenz Huelsbergen, University of Wisconsin-Madison, Computer Science Dept.
Dynamic Program Parallelization

Joseph Hummel, University of California, Irvine, Info. & Computer Science Dept.
Alias Analysis Techniques

Todd Jochem, Carnegie Mellon University, Robotics Institute
Applying Parallel Processing Techniques to Autonomous Vehicle Navigation

Raymond Loy, Rensselaer Polytechnic Institute, Dept. of Computer Science
Parallel Adaptive Strategies for Three-Dimensional Compressible Flows

Catherine McCann, University of Washington, Computer & Science Engineering
Multiprogramming Distributed Memory Systems

Erich Nahum, University of Massachusetts, Computer Science Dept.
Operating System Support for High-Speed Networks